The Privatization of the Internet's Backbone Network

Rajiv C. Shah and Jay P. Kesan

Scholars have neglected the privatization of the Internet's backbone network, despite the obvious significance of the U.S. government turning control of a powerful new communication technology over to the private sector. This article analyzes the transition from a government-sponsored backbone network to multiple commercially owned backbone networks. The authors also analyze the implications of the privatization on the Internet's governance, competition, and performance.

Histories of the Internet abound (Abbate, 1999; Hauben & Hauben, 1997; Kahn & Cerf, 1999; Mowery & Simcoe, 2002; Mueller, 2002; Naughton, 2000; Salus, 1995; Schiller, 1999), yet a comprehensive account of the privatization of the Internet's backbone network does not exist. When it comes to describing the transition of control from the government to the private sector, the descriptions suddenly shift to the passive voice. It becomes unclear who the actors were and what actions they took to privatize the backbone network. The lack of coverage from the mainstream press led Project Censored to place the privatization of the Internet backbone in their top 10 list for 1995 (Jensen, 1997).

This article provides a history of the privatization of the backbone network. This is valuable because the privatization has not been scrutinized by either academics or the press. The privatization is significant for communication scholars for several additional reasons. First, it represents the transfer of an important communication technology to the private sector. The U.S. government spent approximately $160 million in direct subsidies over an 8-year period to fund the backbone network. However, the government likely spent 10 times this amount in developing the Internet through

Rajiv C. Shah *(Ph.D., University of Illinois, Urbana-Champaign; J.D., University of Illinois, Urbana-Champaign) is an Adjunct Assistant Professor in the Department of Communication at the University of Illinois at Chicago. His research interests are in the area of history and public policy implications of communication technology.*

Jay P. Kesan *(Ph.D., University of Texas at Austin; J.D., Georgetown University) is a Professor in the College of Law and the Department of Electrical & Computer Engineering at the University of Illinois at Urbana-Champaign. His research interests are in the area of Internet law and regulation, and intellectual property.*

This research was supported by National Science Foundation Grants ITR-0081426 and IIS-0429217. Opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

 Journal of Broadcasting & Electronic Media 51(1), 2007, pp. 93–109

other public funds from state governments, state-supported universities, and the national government (MacKie-Mason & Varian, 1994). Eventually, the government transitioned the backbone network to private control. Unlike other communication technologies, the privatization of the backbone left few regulatory requirements. As is later pointed out, this has implications for concentration in the backbone industry as well as the government's ability to regulate the Internet. Second, a historical understanding of the privatization provides key insights into contemporary issues, such as the performance of the Internet, competition in the backbone industry, and the governance of the Internet. Finally, this is not the only privatization involving the Internet. The government is now transitioning the Domain Name System (DNS) and the Internet Protocol (IP) address system, both key components of the Internet, to private control. Therefore, any lessons learned from the privatization of the backbone network may aid these other Internet privatizations.

The privatization of the backbone network involved reshaping the National Science Foundation Network (NSFNET) into what is known today as the Internet. The actual privatization consisted of a shift from contracting out a government-subsidized backbone to relying on the market for backbone services. The first part of this article discusses this privatization process. The second part focuses on several important implications of the privatization. Also discussed is how the government lost an opportunity to ensure that societal values were considered in the design of the Internet.

To gain a thorough insight into the privatization process, the authors attempted to gather all known relevant primary and secondary sources. This was accomplished by consulting all known histories of the Internet, Congressional hearings regarding the NSFNET, press accounts, and the landmark compilation of primary documents, known as the *Information Infrastructure Sourcebook* (Kahin, 1993). To corroborate these sources and to provide additional historical depth, the archives of the com-priv mailing group were reviewed. This mailing group was devoted to discussing the privatization process and included leaders from both the government and private sector. Finally, several interviews were conducted to clarify issues raised by the primary and secondary sources.

The Privatization Process

The critical step in the creation of the Internet was the development of the NSFNET. The National Science Foundation (NSF) created this network to link their five supercomputing centers to a wider research community (J. D. Rogers, 1998). By 1986, the NSFNET was operational. It connected to regional networks, which in turn connected to smaller local networks such as universities. The structure of the NSFNET in connecting to the regional networks and the local networks is shown in Figure 1. Eventually the NSFNET connected more than 200 colleges and universities as well as other federal networks, such as the Department of Energy's Energy Sciences Network

Figure 1
Architecture of the NSFNET

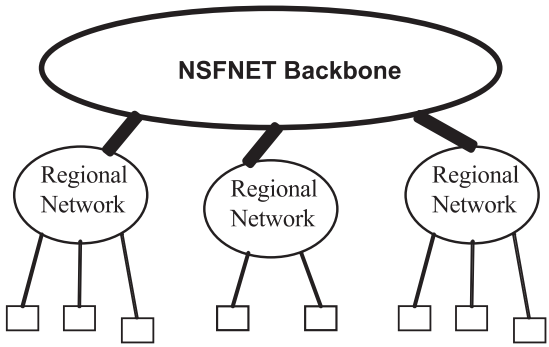

Local Networks, for example, Universities

and NASA's NASA Science Internet (Office of Technology Assessment, 1993). The connection of all these networks via the NSFNET was known as the Internet.

The NSF initially contracted out the management and the operation of the NSFNET backbone to the team of MERIT, a consortium of Michigan universities, in conjunction with IBM and MCI. MERIT was responsible for the management of the project, IBM provided networking equipment, and MCI provided transmission circuits for the network. At the time, the NSFNET was supported by funds from the NSF, state, university, federal, and private sector (*Computer Networks and High Performance Computing,* 1988, pp. 66–67). The remainder of this section describes how the NSF went from subsidizing the backbone network to relying on the market for backbone services. This privatization process is analyzed in three stages. The first concerns the overarching consensus by the government to privatize the NSFNET. The second stage focuses on how the NSF proceeded to privatize the NSFNET. The last stage focuses on the actual privatization plan and redesign of the NSFNET backbone.

Privatization Plan by the U.S. Government

The emphasis on private ownership of telecommunications has a long history in the United States. In 1840, Samuel Morse received funding from the government to demonstrate a radical new communication technology, the telegraph. The importance of this technology led to calls for nationalizing the telegraph. However, Congress did not opt for nationalization and instead permitted private control over telegraph networks (Brock, 1981).

Little has changed 150 years later. The United States continues to have a history of privately owned telecommunication networks (Schiller & Fregoso, 1991). By the late 1980s, both political parties based their telecommunications policy on notions of

deregulation and competition (Olufs, 1999). Experts in computer networks shared this view, whether they were telecommunication executives or academics. Leonard Kleinrock, a prominent computer scientist, told Congress in 1988 that any government-run network would eventually be transitioned to the telecommunications industry (*Computer Networks and High Performance Computing,* 1988). Similarly, Senator Al Gore introduced a bill in 1989 that acknowledged eventual control and ownership of the network by commercial providers (National High Performance Computer Technology Act, 1989).

There were no alternative proposals for privatizing the Internet, such as nationalizing the Internet. According to Lewis Branscomb (1992), Director of the Science, Technology, and Public Policy Program at Harvard University's John F. Kennedy School of Government, individuals and groups associated with planning the NSFNET "take it as a given that commercially provided facilities will be used for the network" (p. 26). In fact, the acceptance of private control was evident in Senator Daniel Inouye's proposal for reserving part of the Internet as a noncommercial space. His proposed bill, the National Public Telecommunications Act of 1994, Senate Bill 2195, did not seek government ownership of the network. He merely wanted to set aside 20% of the network for noncommercial use. However, his proposal met with substantial criticism and never advanced.

General guidelines for privatizing the NSFNET were developed in 1989 by the Federal Coordinating Council on Science, Engineering and Technology (FCCSET; Office of Science and Technology Policy, 1989). The plan called for a three-stage development process for the network. The first stage, which was already underway, was to upgrade networks to T1 (1.5 Mbs/sec) speed. In the second stage, the backbone was to be upgraded, while also providing "a base from which commercial providers can offer compatible networking services nationally" (Office of Science and Technology Policy, 1989, p. 35). The third stage was not expected until the middle or late 1990s and also involved upgrading the backbone. More important, the third stage was to "include a specific, structured process resulting in transition of the network from a government operation to a commercial service" (Office of Science and Technology Policy, 1989, p. 35). The FCCSET plan did not provide details for the privatization process. It did state that the NSF would serve as the lead agency for managing and coordinating the government's backbone network.

NSF Is in Charge

The NSF was responsible for the NSFNET. Although it had contracted with MERIT and its partners to manage the network, the NSF was still responsible for determining how the network was to be used. One critical issue concerned commercial use of the publicly funded NSFNET. The NSF's initial position as stated in its Acceptable Use Policy (AUP) prohibited the use of the NSFNET for purposes not in support of research and education (Kahin, 1992). This policy was consistent with the NSF's mission; the

NSF reasoned that taxpayers would not want a government-subsidized network used for commercial purposes. However, it did relent in a few circumstances. For example, Internet pioneer Vinton Cerf, working for MCI, pushed the government to allow MCI Mail to access federal networks. The government allowed MCI and other providers limited access for "experimental use." The rationale was that these linkages would enhance research and educational uses of the NSFNET by allowing researchers to communicate with more people (Kahin, 1990).

Another purpose behind the AUP was to encourage the growth of commercial backbone companies. After all, the AUP forced commercial customers seeking Internet connectivity to rely on a commercial provider. The NSF believed that the long-term health of computer networking research depended on the growth of these commercial backbone providers. The NSF's logic is summed up by Robert Kahn's aphorism, "You gotta have the flow to get the skim" (Wolff, 1991), the idea being that if the flow or network connectivity becomes exceedingly large and widespread, then it will be possible for the NSF to have a little skim for publicly subsidized computing. This approach was successful as it fostered the creation of commercial backbone providers, most of which were spin-offs from the nonprofit regional networks funded by the NSF (Cerf, 2000). For example, a few of the founders of the nonprofit New York regional network, NSYERNet, established a for-profit company, Performance Systems International (PSI).

The growing number of commercial providers and the push for commercial use led the NSF to develop a plan for transitioning the NSFNET backbone to commercial providers. The Director of the NSF Division of Networking and Communications Research and Infrastructure, Stephen Wolff, led the NSF's efforts to privatize the backbone network. To explore the issues of privatization, Wolff held an invitation-only meeting at Harvard University in March 1990. The 30 people in attendance included economists, public policy specialists, industry representatives, and university networking personnel. A summary report of this meeting was published by Kahin (1990) for the Internet community. Throughout this meeting and later meetings, the NSF's goal was to decide how best to build the market for commercial backbone services, so that the government-subsidized NSFNET would no longer be needed. By late 1991, after several more workshops and meetings, the NSF released a Project Development Plan (National Science Foundation, 1992). It envisioned the creation of a new network structure supported by multiple commercial backbone providers within 18 months (Wolff, 1992).

While the NSF was planning the privatization process, its contractors unexpectedly moved forward in selling connectivity to the NSFNET by creating a not-for-profit corporation, Advanced Network Services (ANS), in late 1990 ("MCI and IBM Form Nonprofit Supercomputing Company," 1990). ANS entered into a subcontracting agreement with MERIT, IBM, and MCI to manage the NSFNET. The NSF permitted these actions. ANS next proceeded to set up a for-profit subsidiary that offered commercial networking services on the same infrastructure as the NSFNET (Valovic, 1991). William Schrader of PSI viewed this as unfair competition because his firm could not pro-

vide connectivity across the NSFNET (Markoff, 1991). However, the NSF had agreed to allow ANS to solicit commercial customers, although the NSF failed to publicly announce this decision. The complaints over ANS led to Congressional hearings on the management of the NSFNET in March 1992 (*Management of NSFNET,* 1992).

At the hearings, competitors of ANS described the uneven playing field. Schrader stated the NSF's actions put ANS in a "monopoly position" and was akin to "giving a federal park to K-mart" (*Management of NSFNET,* 1992, pp. 87–88). Mitch Kapor of the Electronic Frontier Foundation and the Commercial Internet Exchange favored more constructive criticism by stating "recent negative publicity surrounding the NSFNET has come because important decisions about the network were made without opportunity for public comment or input from commercial Internet providers" (*Management of NSFNET,* 1992, p. 85). In response, Congressman Boucher asked the Office of the Inspector General (OIG) to review the NSFNET's management (OIG, 1993).

The OIG report concluded that the NSF did nothing illegal. Nevertheless, the report contained strongly worded recommendations for the NSF. The OIG report found that a key source of the problems was that the NSF had not anticipated commercial use of the NSFNET in its original agreement with MERIT and its partners. This allowed ANS to take advantage of the situation. The OIG report noted that the NSF did not act as a government agency should act. Although the NSF's actions were legal, the NSF did not adequately document and make public its relationship and agreement with ANS. The OIG report sent a clear message to the NSF—it needed to be more transparent and receptive to other commercial backbone companies to ensure a level playing field. The NSF would respond positively and address these concerns during the redesign of the NSFNET. For example, the NSF's transition process would ensure that all commercial backbone providers could equally participate in the privatized network. However, the NSF's perceived partiality to ANS served to confirm doubts for the commercial backbone providers that government was incapable of supervising a commercial backbone system.

Congress still allowed the NSF to continue the privatization process. However, it was clear that the NSF's actions were going to be carefully scrutinized. Congress also responded to the concerns among commercial backbone providers over the limitations found in the AUP. A law was passed allowing commercial use of the NSFNET as long as it would increase the networks' utility for research and education ("Scientific and Advanced-Technology Act," 1992). This new, liberal AUP permitted further commercial use and led to more growth for the NSFNET.

NSF's Final Plan for the NSFNET

In June 1992, the NSF released a draft solicitation for public comment on its proposed changes to the NSFNET. The new design was developed internally by three of its engineers (Aiken, Braun, Ford, & Claffy, 1993). In response, the NSF received more

than 240 pages of comments, the majority coming from industry groups. The comments largely expressed concern that only a few firms would be awarded the entire contract and, suggested instead, the desire for a more competitive network design. Several of these comments also noted potential long-term problems with the redesign. The NSF never addressed these comments and these issues were left to the backbone providers. The next section focuses on the implications of this inaction.

In the spring of 1993, the NSF released a revised solicitation in response to the comments on the proposed changes to the NSFNET (NSF, 1993). The new design had three parts. The first part consisted of a Routing Arbiter, which operates as a "traffic cop" to ensure consistent routing policies. The second part proposed the creation of a Very High Speed Backbone Service (vBNS) to replace the NSFNET as a new high-speed backbone for research and educational use. The third part of the revised solicitation concerned the use of network access points (NAPs) to connect together the vBNS, federal networks, and commercial backbone networks. The concept of multiple backbones interconnecting through NAPs is shown in Figure 2. Instead of a central backbone connecting the regional networks, the regional networks had to choose a commercial backbone network. This network structure ensured a place for multiple commercial backbone providers. To ensure interconnectivity, the commercial backbone networks would interconnect at the NAPs. This new design appeared to provide a level playing field for commercial backbone providers, because any backbone could connect at the NAPs.

In February 1994, the NSF awarded contracts to establish four NAPs. The contracts were to Sprint for New York, MFS for Washington, DC, Ameritech for Chicago, and

Figure 2
The New Design for the NSFNET

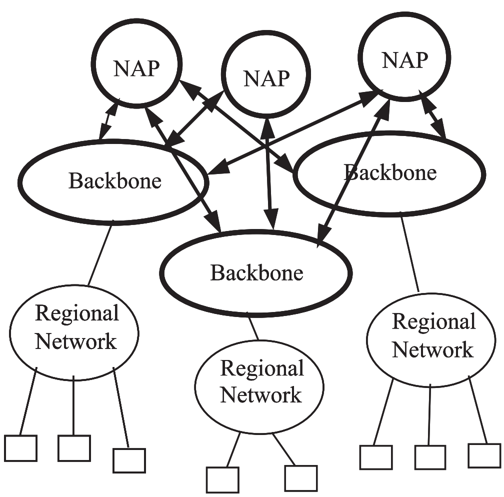

Pacific Bell for California. By October, the regional networks were supposed to purchase connectivity from a commercial provider and disconnect from the NSFNET backbone. None of the regional networks were able to meet the October deadline, but 6 months later they were migrated to commercial providers (Fazio, 1995). On April 30, 1995, the NSFNET was retired. The government had successfully transitioned the network to multiple commercial backbone providers.

A year later, the NSF ended its contracts for the four public NAPs. In effect, this transferred control of the NAPs to their private-sector contractors (Farnon & Huddle, 1997). By this point, the government had transitioned from contracting out services to allowing the market to provide Internet backbone service—thus fully privatizing the Internet backbone network. The government had no ongoing role in overseeing or regulating the backbone network. After this decision, the only remaining NSF-subsidized backbone service was the vBNS. The vBNS served as the initial foundation of the ongoing Internet2 networking effort.

Implications of the Privatization

The privatization of the NSFNET backbone fundamentally reshaped the infrastructure of the Internet. Instead of one major backbone—the NSFNET—the new network depended on multiple backbone providers. The NSF created the NAPs to interconnect these networks and to prevent a balkanized Internet. This new network was designed to create competition for backbone services. In designing this new network, the NSF put few constraints on the design and use of the network. After all, commercial backbone providers had no taste for any continuing governmental role.

This section discusses how the NSF's hands-off approach shaped the Internet with predictable implications. First, the lack of any performance requirements for the public NAPs led to them becoming congested. This fostered new technological innovations, which affect how information is transferred across the Internet. Second, the lack of an interconnection policy has resulted in a concentrated industry with a few large backbone providers. As a result, there is a high entry barrier to the backbone market along with concerns that a few backbone providers may act anticompetitively. Third, the government's hands-off regulatory policy ensures that private industry governs and controls the Internet's backbone. A final implication concerns the lost opportunity by government to ensure that societal values were considered in the transition to multiple commercial backbone networks.

Performance of the Network Access Points

One of the issues raised during the NSF's redesign was placing performance requirements for the NAPs. In designing the NAPs, the NSF neglected to put into place minimum performance requirements. Instead, the NSF relied on the NAP operators to

ensure adequate performance. Such a policy resonates with the NSF's deference to commercial backbone providers. The NAPs were not able to maintain their performance as more networks connected to them, and as a result, they soon became congested. For example, Sprint's backbone network operates at 2.5 billion bits per second, but the highest speed offered at the Sprint-managed public NAP is 45 million bits per second. This is aptly characterized as "giving drivers on a six-lane highway access via a dirt road" (Weinberg, 2000, p. 238). Not surprisingly, many backbone providers consider the public NAPs' connections worthless.

The congestion has led to the use of private exchange points (Winkleman, 1998). By interconnecting at private exchange points, it is possible to avoid routing traffic through the public NAPs. For example, the five largest backbone providers have each implemented at least four connections with the other large backbone providers at private exchanges. This allows the large backbone providers to provide high-performance Internet access for their customers. Meanwhile, smaller networks are often relegated to the congested public NAPs, because there is no requirement for a private exchange to treat other networks in a fair and nondiscriminatory manner (Caruso, 2000).

The congestion at the NAPs has also led to the use of new technologies, which fundamentally affects how information is transferred across the Internet. The first is the content delivery networks (CDNs) such as Akamai, which attempt to bypass as much of the public Internet as possible. Instead of having one server for a Web site's content, a CDN uses a number of servers with replicated content at the edge of the network. Users are then transparently routed to the nearest available server. This allows for content to be delivered much more quickly to users. This service is not free. As a result, Web sites that can afford to pay for a CDN will be able to provide richer content much more quickly than other Web sites (Sandvig, 2006).

A second technology used to address congestion is known as Quality of Service (Ferguson & Huston, 1998). This technology modifies one of the basic tenets of Internet traffic. Traditionally, all Internet traffic was treated alike and everyone would give their best effort to transmit all traffic. Quality of Service provides preferential treatment to traffic. The ability of service providers to segregate traffic by application, users, or network can have enormous implication for society's use of the Internet. One controversial use of Quality of Service is by cable companies that expedite the delivery of affiliated content while demoting competitive material to second-class network service (Lemley & Lessig, 2001).

The use of Quality of Service favors large backbone providers over smaller ones, because effective Quality of Service relies on maintaining control of the network traffic over the entire length of the network. Under the current system, smaller networks depend on a few large backbone providers for some of their Internet transport. This limits their use of Quality of Service and provides the large backbone providers with a technological edge. This is an example of a long-term consequence of the lack of performance requirements at the NAPs.

Competition and the Lack of an Interconnection Policy

The government has historically used interconnection policies to require common carriers to maintain a duty to serve and interconnect on a nondiscriminatory basis. For example, steamships, railroads, and the telephone network were all subject to policies to ensure interconnection (Speta, 2002). Similarly, the Telecommunications Act of 1996 requires telecommunications carriers to interconnect, because Congress implicitly understood that a telecommunication network's value increased as more people became connected to the network. Although these obligations and laws apply to traditional voice telecommunication carriers, the Federal Communications Commission (FCC) has held that Internet backbone providers are not telecommunications carriers as defined in the 1934 or 1996 Telecommunications Acts (Kende, 2000).

During the development of the telephone, AT&T often refused to interconnect with its rivals. This may have contributed to the initial growth of local telephone infrastructure (Mueller, 1997), but it also fostered the development of a monopolistic telephone network. It was through network effects that AT&T was able to leverage its larger long-distance network over smaller competitors in the early 20th century (Brock, 1981). There are similar concerns for the Internet (Cremer, Rey, & Tirole, 2000). This is why the U.S. government forced MCI to sell its Internet assets before the merging with WorldCom in 1998. Regulators were concerned that if MCI WorldCom controlled 70% of the Internet, this level of concentration would greatly impair competition in the backbone industry (FCC, 1998).

The lack of an interconnection policy favors large backbone providers over smaller ones in several ways (Frieden, 1998). First, large backbone providers do not reveal what other networks they peer with or their terms for the peering. This secrecy benefits the large backbone providers as they undoubtedly claim they are well connected. Second, smaller networks often have to sign nondisclosure agreements, which limits their ability to share information on industry interconnection terms. This practice limits the creation of industry standards concerning who should get free peering as well as determining reasonable rates (Weinberg, 2000). Third, larger backbone providers can discontinue their interconnection arrangements with little notice (Cukier, 1997). For example, in 2005, Level 3 Communications decided to no longer exchange traffic with the smaller provider Cogent. This left Cogent's customers only partially connected to the Internet (Mohammed, 2005).

Several commentators have proposed interconnection policies as a method for promoting competition and ensuring the stability of network connections. Speta (2002) argued that Internet backbone providers should be treated as common carriers and, therefore, have to interconnect on a nondiscriminatory basis. Another proposal, by Varian (1998), stems not from common carriers, but from a standard used in intellectual property licensing. He proposed an interconnection policy for backbone providers using a "fair, reasonable, and nondiscriminatory" standard. Varian's preferred method for enforcement is an industry-wide arbitration board. He believes an arbitra-

tion board could better adapt to changes in circumstances and technology than government.

The lack of an interconnection policy has clear implications for competitive entry into the backbone industry, which means increased costs for consumers and businesses. In 1995, there were five major backbone providers—UUNET, ANS, SprintLink, BBN, and MCI—with a combined market share of 90% (Winkleman, 1998). In 2000, there were five major backbone providers: MCI WorldCom (which bought ANS and UUNET), Sprint, Genuity (formerly GTE, which bought BBN), AT&T, and Cable & Wireless (which owns MCI's old backbone; Weinberg, 2000). AT&T built its network with its own large fiber optic network, by acquiring IBM's Global Network, and by acquiring one of the early Internet backbone providers, CERFnet. Despite the rapid growth in Internet services and new Internet start-ups, such as Level 3 Communications and Qwest, the same few companies still control the Internet. After all, the only way to become a large backbone provider is to buy one!

Regulation of Internet Backbone Providers

In today's privatized Internet, the private sector governs the Internet. The FCC's lenient regulatory approach began in the early 1980s when the FCC ruled that enhanced services were not subject to common carrier regulation (Oxman, 1999). This removed Internet services from the many regulations that govern voice telephony. In contrast, many Internet service providers benefit from FCC-regulated interconnection policies with local phone companies.

The prevailing wisdom, as represented by David Farber, former Chief Technologist at the FCC, is that the commercial backbone companies have a self-interest in the success of the Internet (Caruso, 2000). The reliance on self-interest for managing the Internet has worked well in maintaining system performance and reliability. For example, the Internet was able to easily withstand the shutdown of the Ebone network owned by KPNQwest. Traffic was routed onto other networks, which resulted in slight delays for Internet traffic in Europe, but no major problems (Hansell, 2002).

One of the consequences of the lack of regulatory oversight is that government has trouble intervening when there are problems. For example, former FCC Chairman Michael Powell acknowledged that although he could not stop backbone providers from shutting down their service, he could prevent phone service from being stopped. Moreover, backbone providers could legally ignore FCC directives and cut off service to customers with little notice (Dreazen, 2002). This type of rapid shutdown has happened. Cable & Wireless unilaterally terminated the backbone provider PSI. This resulted in customers unexpectedly without Internet service (Kraph, 2001). Similarly, when Excite@Home went bankrupt it left nearly a million cable modem users without Internet access for a week. These interruptions would not be permitted for phone service due to stricter regulatory oversight.

Lost Opportunity by the NSF

During the redesign of the backbone, the NSF was charged with creating guidelines for acceptable conduct as well as technical requirements for the backbone. The norms and technical requirements developed by the NSF are analogous to rules or law for the backbone (Lessig, 1999). During the creation of a competitive backbone structure, the NSF also lost an opportunity to address related societal concerns during the redesign. For example, it was well understood that the NSFNET needed to improve its security (Office of Science and Technology Policy, 1987). In 1991, Charles Brownstein, assistant director in the NSF's Computer and Information Science and Engineering Directorate, acknowledged that security "is a big concern and growing bigger every day" (B. Rogers, 1991). However, during the NSFNET's lifetime, the government never seriously addressed concerns of security. The exception is the CERT Coordination center, which serves as a reporting center for Internet security problems.

The significant point is that government could have done something, not necessarily that it should have. There was little public debate on the role of government. The privatization process was limited to industry insiders and the participating public interest groups only sought competition. As Hauben and Hauben (1997) noted, it was not until November 1994 that the U.S. government attempted to create a dialogue with the Internet community on the privatization process. By then the privatization was a fait accompli.

The NSF could have prepared the Internet for its transition from an academic network to a commercial one. One way this could have been done is by requiring backbone providers using the government-sponsored NAPs to upgrade their technology. For example, the design of the next-generation Internet Protocol, IPv6, incorporates heightened security measures that can encrypt the contents of a message as well as reliably verify the sender of a message (Hagen, 2002). These security measures would have tangible effects, such as limiting junk e-mail. Other improvements of IPv6 include a larger address space, which limits the utility of attacks based on scanning a whole network. Besides addressing security, the NSF could have set guidelines for how the network is to be used. The NSF could have ensured that no network preferentially favors a certain application. The FCC is now considering this network neutrality policy. However, because the NSF never even considered any such actions, one can only wonder what may have resulted if the NSF had used its one-time opportunity to shape the Internet.

Discussion

The privatization of the backbone network was the result of the U.S. government's desire to rely on commercial backbone providers to manage the Internet. This desire was manifested in decisions to allow commercial content across the NSFNET, policies that encouraged the growth of commercial backbone providers, and movement to

privatize the NSFNET. The NSF was given responsibility for the privatization process. However, the NSF was not fully prepared for the privatization process. In retrospect, it is not surprising that there were these types of problems. The NSF's mission is to manage grants for basic science, not to supervise the privatization of government resources. Inevitability, the NSF ran into problems and was publicly admonished by Congress.

There are lessons to be learned from the NSF's role in the privatization process. First, the NSF needs to exercise caution when entering into contracts that may result in eventual privatization. Many of the problems with the privatization process could have been avoided if the NSF had considered the issue of commercial use in its agreement with MERIT and its partners. Second, maybe the NSF is not the best agency to handle the privatization process. After all, a few years after the privatization of the NSFNET backbone network, the NSF refused to manage the privatization of the DNS. The White House then intervened and assigned responsibility for privatizing the DNS to the Commerce Department (Mueller, 2002). Finally, the NSF's reluctance to place any constraints on the design and use of the network had several predictable results. These consequences could have been avoided, but the NSF chose a hands-off policy that prevented the agency from setting guidelines and standards.

The lack of guidelines and standards led to several predictable implications. The first result concerns the performance issues with the Internet. The NSF never put into place minimum performance requirements for the NAPs, which meant they became congested. This has forced backbone providers to negotiate for access to private exchange points and utilize technologies to manage network traffic. These solutions favor deep-pocketed content providers and also disparately impact the smaller commercial backbone providers.

The second is how the lack of an interconnection policy has led to a concentrated market and barriers to entry for new backbone providers. The lack of an interconnection policy provides large providers with a great deal of power, for example, by allowing them to arbitrarily decide whether to provide connectivity to smaller providers. This uneven competitive environment is exacerbated as the backbone market concentrates as a result of mergers.

Third, the NSF's reluctance to place any limits on the network was followed by policymakers at the FCC. This reluctance overlooks potential conflicts of interest that may exist between large backbone providers and their competitors, customers, or some public interest. Government may need to intervene if a market failure exists. Examples of these conflicts and market failures vary from updating the regulations for Communications Assistance for Law Enforcement Act and interconnections fees to meeting new challenges such as Voice over Internet Protocol (VoIP) and spam. Policymakers within the FCC are now struggling with regulating something previously defined as unable to be regulated. This could have been anticipated. The FCC should have been prepared to consider options for how best to regulate the Internet.

The final implication has to do with lost opportunities. The privatization process could have better designed the Internet for commercial use, for example, for security.

Not using the process in this way has resulted in scrambling to patch the Internet as a cure, instead of considering whether this was preventable. The point is not what the government did not do, but rather what the government never considered.

Conclusion

This article describes the privatization process for the backbone network as well some of the implications of the privatization. This telling of the history identified three important facets of the privatization. First, several reasons for the privatization were identified. These include privatization as a common goal, the desire to send commercial content across the Internet, and transitioning control over the Internet's infrastructure to multiple commercial backbone providers. Second, the primary actor in this process was the NSF. Finally, the privatization was manifested by a shift from a NSF-subsidized central backbone to relying on multiple commercial backbone providers for Internet connectivity.

The authors identified problems with the NSF-managed privatization process. The goal is not to scapegoat the NSF, but to analyze why these problems occurred, so as to learn from them. It is not surprising that the NSF had problems; after all, managing a privatization process is far from their core mission. At several points, the NSF's omissions had predictable consequences for shaping the Internet. Specifically, concerning the Internet's performance, issues of competition and concentration within the backbone industry, and the problems with regulating something that is defined as not capable of being regulated. Finally, the NSF lost an opportunity to shape the commercial Internet to comport with societal goals, such as security.

These implications are valuable, because they provide insights into why the Internet is the way it is today. Moreover, these implications are also relevant to the other privatizations of the Internet consisting of the DNS and the allocation of IP addresses. It is well established that the privatization of the DNS is already suffering from problems with competition and governance (Froomkin, 2000; Kesan & Shah, 2001; Mueller, 2002). The hope is that this work will aid in preventing these problems from being repeated in future privatizations.

References

Abbate, J. (1999). *Inventing the Internet.* Cambridge, MA: MIT Press.

Aiken, R., Braun, H.-W., Ford, P., & Claffy, K. (1993). NSF implementation plan for interagency interim NREN. *Journal of High Speed Networks, 2*(1), 1–25.

Branscomb, L. M. (1992). Information infrastructure for the '90s: Public policy perspective. In B. Kahin (Ed.), *Building information infrastructure* (pp. 15–30). New York: McGraw-Hill Primis.

Brock, G. W. (1981). *The telecommunications industry: The dynamics of market structure.* Cambridge, MA: Harvard University Press.

Caruso, D. (2000, February 14). The Internet relies on networks' passing data to one another. But what happens if one of them refuses? *New York Times,* p. C4.

Cerf, V. G. (2000). Computer networking: Global infrastructure for the 21st century. Retrieved September 10, 2006, from http://www.cs.washington.edu/homes/lazowska/cra/networks.html

Computer networks and high performance computing: Hearings before the Subcommittee on Science, Technology, and Space of the Senate Committee on Commerce, Science, and Transportation, 100th Cong. (1988).

Cremer, J., Rey, P., & Tirole, J. (2000). Connectivity in the commercial Internet. *Journal of Industrial Economics, 48,* 433–472.

Cukier, K. N. (1997, December). *Peering and fearing: ISP interconnection and regulatory issues.* Paper presented at the Impact of the Internet on Communications Policy conference, Cambridge, MA.

Dreazen, Y. (2002, July 31). FCC chief cites Internet weak spot. *Wall Street Journal,* p. A2.

Farnon, M., & Huddle, S. (1997). Settlement systems for the Internet. In B. Kahin & J. H. Keller (Eds.), *Coordinating the Internet* (pp. 377–403). Cambridge, MA: MIT Press.

Fazio, D. (1995, March 14). Hang onto your packets: The Information Superhighway heads to Valleyfair. Retrieved September 10, 2006, from http://ei.cs.vt.edu/~history/internet.hist.html

Federal Communications Commission. (1998). *Application of WorldCom, Inc. and MCI Communications Corporation for transfer of control of MCI Communications Corporation to WorldCom, Inc.* (No. CC Docket No. 97–211).

Ferguson, P., & Huston, G. (1998). *Quality of service: Delivering QoS on the Internet and in corporate networks.* New York: Wiley.

Frieden, R. (1998). Without public peer: The potential regulatory and universal service consequences of Internet balkanization. *Virginia Journal of Law and Technology, 3,* 8.

Froomkin, M. (2000). Wrong turn in cyberspace: Using ICANN to route around the APA and the Constitution. *Duke Law Journal, 50,* 17–186.

Hagen, S. (2002). *IPv6 essentials.* Sebastopol, CA: O'Reilly & Associates.

Hansell, S. (2002, July 24). Among WorldCom puzzles, future of UUNet service*. New York Times,* p. C.8.

Hauben, M., & Hauben, R. (1997). *Netizens: On the history and impact of the Usenet and the Internet.* Los Alamitos, CA: IEEE Computer Society Press.

Jensen, C. (1997). *20 years of censored news.* New York: Seven Stories Press.

Kahin, B. (1990, November). RFC 1192: Commercialization of the Internet. November. Retrieved September 10, 2006, from http://www.faqs.org/rfcs/rfc1192.html

Kahin, B. (1992). The NREN as information market: Dynamics of public, private, and academic publishing. In B. Kahin (Ed.), *Building information infrastructure* (pp. 323–401). New York: McGraw-Hill Primis.

Kahin, B. (Ed.). (1993). *Information infrastructure sourcebook.* Cambridge, MA: John F. Kennedy School of Government & Harvard Office for Information Technology.

Kahn, R. E., & Cerf, V. G. (1999, December). What is the Internet (and what makes it work). Retrieved September 10, 2006, from http://www.policyscience.net/cerf.pdf

Kende, M. (2000). *The digital handshake: Connecting Internet backbones* (No. OPP Working Paper No. 32). Washington, DC: Federal Communications Commission, Office of Plans and Policy.

Kesan, J. P., & Shah, R. C. (2001). Fool us once shame on you—Fool us twice shame on us: What we can learn from the privatizations of the Internet backbone network and the Domain Name System. *Washington University Law Quarterly, 79,* 89–220.

Kraph, E. (2001, July). Straightening out the Internet backbone. *Business Communication Review, 31,* 14–15.

Lemley, M. A., & Lessig, L. (2001). The end of end-to-end: Preserving the architecture of the Internet in the broadband era. *UCLA Law Review, 48,* 925–972.

Lessig, L. (1999). *Code and other laws of cyberspace.* New York: Basic Books.

MacKie-Mason, J. K., & Varian, H. (1994). The economics of the Internet. *Dr. Dobb's Journal, 19*(14), 6–16.

Management of NSFNET: Hearings before the Subcommittee on Science of the House Committee on Science, Space, and Technology, 102d Cong. (1992).

Markoff, J. (1991, December 19). Data network raises monopoly fear. *The New York Times,* p. D.7.

MCI and IBM form nonprofit supercomputing company. (1990, September 18). *Communications Daily, 10,* 3.

Mohammed, A. (2005, October 14). Internet access dispute cut off some businesses. *Washington Post,* p. D4.

Mowery, D. C., & Simcoe, T. (2002). Is the Internet a U.S. invention? An economic and technological history of computer networking. *Research Policy, 31,* 1369–1387.

Mueller, M. L. (1997). *Universal service: Competition, interconnection and monopoly in the making of the American phone system.* Washington, DC: AEI.

Mueller, M. L. (2002). *Ruling the root: Internet governance and the taming of cyberspace.* Cambridge, MA: MIT Press.

National High Performance Computer Technology Act, S. 1067, 101st Cong. (1989).

National Science Foundation. (1992, June). Public Draft: Network Access Point Manager/Routing Authority and Very High Speed Backbone Network Services provider for NSFNET and the NREN program. Retrieved September 10, 2006, from http://www.eff.org/Infrastructure/Govt_docs/nsf_nap_vbns_original_92.rfc

National Science Foundation. (1993, May). Network access point manager, routing arbiter, regional network providers, and very high speed backbone network services provider for NSFNET and the NREN program. Retrieved September 10, 2006, from http://www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf9352

Naughton, J. (2000). *A brief history of the future: From radio days to Internet years in a lifetime.* Woodstock, NY: Overlook.

Office of Inspector General. (1993). *OIG review of NSFNET* (No. OIG9301). Washington, DC: National Science Foundation.

Office of Science and Technology Policy. (1987). *A research and development strategy for high performance computing.* Washington, DC: Executive Office of the President.

Office of Science and Technology Policy. (1989). *The federal high performance computing program.* Washington, DC: Executive Office of the President.

Office of Technology Assessment. (1993). *Advanced network technology—Background paper* (No. OTA-BP-TCT-101). Washington, DC: U.S. Government Printing Office.

Olufs, D. (1999). *The making of telecommunications policy.* Boulder, CO: Lynne Rienner.

Oxman, J. (1999). *The FCC and the unregulation of the Internet* (No. OPP Working Paper No. 31). Washington, DC: Office of Plans and Policy, Federal Communications Commission.

Rogers, B. (1991, September 30). Network's goal to push the edge of technology: NSF helps keep rapidly growing Internet on the cutting edge of research. *Government Computer News, 10,* 49.

Rogers, J. D. (1998). Internetworking and the politics of science: NSFNET in Internet history. *The Information Society, 14,* 213–228.

Salus, P. H. (1995). *Casting the Net: From ARPANET to Internet and beyond.* Reading, MA: Addison-Wesley.

Sandvig, C. (2006). The structural problems of the Internet for cultural policy. In *Critical cyberculture studies: Current terrains, future directions.* New York: New York University Press, 173–190.

Schiller, D. (1999). *Digital capitalism: Networking the global market system.* Cambridge, MA: MIT Press.

Schiller, D., & Fregoso, R. (1991). A private view of the digital world. *Telecommunications Policy, 15,* 195–208.

Scientific and Advanced-Technology Act, 42 U.S.C. § 1862(g) (1992).

Speta, J. B. (2002). A common carrier approach to Internet interconnection. *Federal Communications Law Journal, 54,* 225–280.

Valovic, T. (1991, July). ANS and Internet commercialization. *Telecommunications Americas, 25,* 4.

Varian, H. R. (1998, June 8). How to strengthen the Internet's backbone. *Wall Street Journal,* p. A22.

Weinberg, N. (2000, June 12). Backbone bullies. *Forbes,* pp. 236–238.

Winkleman, J. (1998, August 15). Getting connected. *America's Network, 102,* 30–34.

Wolff, S. (1991, May 14). NREN. Message posted to com-priv electronic mailing list, archived at http://diswww.mit.edu/menelaus/com-priv/686

Wolff, S. (1992, November 26). NSFNET backbone services after November 1992. Retrieved March 31, 1999, from ftp://nic.merit.edu/cise/pdp.txt